Online Dual Coordinate Ascent Learning

Bicheng Ying, Kun Yuan, and Ali H. Sayed

Department of Electrical Engineering
University of California, Los Angeles

arXiv:1602.07630v1 [math.OC] 24 Feb 2016

***Abstract*—The stochastic dual coordinate-ascent (S-DCA) technique is a useful alternative to the traditional stochastic gradient-descent algorithm for solving large-scale optimization problems due to its scalability to large data sets and strong theoretical guarantees. However, the available S-DCA formulation is limited to finite sample sizes and relies on performing multiple passes over the same data. This formulation is not well-suited for online implementations where data keep streaming in. In this work, we develop an *online* dual coordinate-ascent (O-DCA) algorithm that is able to respond to streaming data and does not need to revisit the past data. This feature embeds the resulting construction with continuous adaptation, learning, and tracking abilities, which are particularly attractive for online learning scenarios.**

***Index Terms*—Online algorithm, dual coordinate-ascent, stochastic gradient-descent, stochastic proximal gradient, adaptation, learning, support-vector machine.**

I. Introduction and Related Work

We consider minimizing a regularized stochastic convex risk function of the form:

$$\min_{w} J(w) \;\triangleq\; \mathbb{E}\, Q(w;\boldsymbol{x}) + \rho R(w) \tag{1}$$

where the expectation is over the distribution of the data represented by the boldface letter $\boldsymbol{x}$, $w \in \mathbb{R}^M$ is an unknown parameter vector, and $\rho \geq 0$ is a scaling factor. Moreover, the loss $Q(w;\boldsymbol{x})$ is a convex function over w and it may be non-differentiable. The term $R(w)$ is a strongly-convex regularization factor such as ℓ_2 or elastic-net regularization. In learning applications, it is customary for the data to consist of a scalar variable $\boldsymbol{\gamma}$ and an $M-$dimensional feature vector, $\boldsymbol{h}$, i.e., $\boldsymbol{x} = \{\boldsymbol{\gamma}, \boldsymbol{h}\}$. We assume that the loss function depends on the data in the following manner:

$$Q(w;\boldsymbol{x}) \;\triangleq\; Q(\boldsymbol{h}^{\mathsf{T}} w; \boldsymbol{\gamma}) \tag{2}$$

This problem formulation is typical of many scenarios including leaky least-mean-squares [1], support-vector machines [2], [3], regularized logistic regression [4], [5], and others.

In practice, it is customary to replace problem (1) by the minimization of a regularized *empirical* risk that is based on a collection of N data points, namely,

$$w^o \;\triangleq\; \arg\min_{w} \;\; \frac{1}{N}\sum_{n=1}^{N} Q\left(h_n^{\mathsf{T}} w; \gamma(n)\right) + \rho R(w) \tag{3}$$

Here, the data $\{\gamma(n), h_n\}$ represent realizations arising from the distribution driving the variables $\{\boldsymbol{\gamma}, \boldsymbol{h}\}$ and N is the size of the data sample. One traditional, yet powerful, approach to solving problem (3) is to employ a stochastic (sub)gradient method (SGD, for short) [1], [6]–[8]. This method can be implemented both in empirical form (involving repeated passes over a finite data sample) or online form (in response to streaming data). Either way, when a constant step-size, μ, is used to drive the iteration, the SGD method can be shown to converge exponentially fast to a small neighborhood around w^o with a steady-state error variance that is on the order of $O(\mu)$. Therefore, sufficiently small step-sizes can ensure satisfactory steady-state performance albeit at the cost of a slower convergence rate [5], [9], [10].

An alternative approach is to solve problem (3) in the dual domain. Instead of minimizing (3) directly, one can maximize the dual cost function using a coordinate-ascent algorithm [11]. The dual problem involves maximizing over N dual variables. Since updating all N dual variables at each iteration can be costly, the coordinate-ascent implementation updates one dual variable at a time. There have been several recent investigations along these lines in the literature with encouraging results. For example, references [3], [12] observed that a dual coordinate-ascent (DCA, for short) method can outperform the SGD algorithm when applied to large-scale SVM. Later, a *stochastic* version of DCA (denoted by S-DCA) was examined in [13], [14] for more general risk functions. Compared with DCA, at each iteration, the stochastic implementation picks one data sample randomly (not cyclically) and updates the corresponding coordinate. Reference [13] showed that S-DCA converges exponentially to the exact minimizer w^o by running repeated passes over the finite data sample, which is a notable advantage over SGD.

Despite the apparent advantages in terms of theoretical guarantees and experimental performance, the stochastic DCA implementation suffers from three drawbacks for online scenarios. First, the available S-DCA implementation needs to know beforehand the size of the training data, N, since this value is explicitly employed in the algorithm. When data streams in, the value of N is constantly changing and, therefore, the S-DCA implementation will not be applicable. Second, S-DCA needs to perform multiple passes over the same finite data sample. This situation is problematic for online operation when new data keeps streaming in and it is not practical to keep examining past data. Third, the S-DCA algorithm assigns the same weight $(1/N)$ to each data sample in the training set. This is not ideal for scenarios where the minimizer w^o can drift with time since it deprives the algorithm of adaptation and tracking abilities.

In summary, while stochastic (sub)gradient techniques are

This work was supported in part by NSF grants CCF-1524250 and ECCS-1407712, and by DARPA project N66001-14-2-4029. Emails:{ybc,kunyuan,sayed}@ucla.edu

able to solve problems of the type (3) in an online manner, the available DCA and S-DCA algorithms lose this important feature. Motivated by these considerations, we focus in this article on developing an *online* stochastic coordinate-ascent algorithm, denoted by the letters O-DCA. While it shares some useful features with S-DCA, the online version allows the sample size to increase and continuously adjusts the weights that are assigned to the samples. The experimental results in this work illustrate the superior performance of O-DCA over SGD in terms of convergence rate and accuracy. We comment on these results by explaining how O-DCA shares interesting and revealing connections with stochastic gradient-descent and stochastic proximal gradient algorithms in the primal domain. Specifically, we will show that under ℓ_2-regularization, the proposed O-DCA algorithm is related to a stochastic proximal gradient implementation, which helps explain the observed superior performance of O-DCA over SGD.

II. Problem and Algorithm Formulation

A. Dual problem

We first replace the empirical problem (3) by a more general weighted formulation that is able to capture several scenarios of interest as special cases. Namely, we consider instead the following problem:

$$\min_{w} \quad \frac{1}{\Delta_N}\sum_{n=1}^{N}\delta_{n,N}Q\left(h_n^{\mathsf{T}}w;\gamma(n)\right)+\rho R(w) \tag{4}$$

where $\delta_{n,N}\geq 0$ is a weighting scalar factor and $\Delta_N>0$ is a normalization scalar factor. Both factors depend on the number N of data points. Different choices for these factors correspond to different useful situations [1], [15]:

(C1) (**Infinite-length window**): This case corresponds to the choice $\delta_{n,N}=1$ and $\Delta_N=N$, which reduces to (3). For these choices, all data starting from the remote past are scaled similarly.

(C2) (**Exponential-weighting window**): In this case, we set

$$\delta_{n,N}=\beta^{N-n},\quad \Delta_N=\sum_{n=1}^{N}\beta^{N-n}=\frac{1-\beta^N}{1-\beta} \tag{5}$$

for some forgetting factor $\beta\in(0,1)$. Usually, the value of β is very close to one, so that recent data are weighted more heavily than data from the remote past.

(C3) (**Finite-length sliding window**): In this case, we focus on the most recent L data points by setting $\Delta_N=L$ (for the initial stages when $N\leq L$, we set $\Delta_N=N$) and

$$\delta_{n,N}=\begin{cases}1, & \text{when } n>N-L\\ 0, & \text{when } n\leq N-L\end{cases} \tag{6}$$

In order to examine the dual problem of (4), we first rewrite it in the following equivalent form involving a set of linear constraints in terms of scalar variables $\{z(n)\}$:

$$\min_{w,\{z(n)\}} \quad \frac{1}{\Delta_N}\sum_{n=1}^{N}\delta_{n,N}Q\left(z(n);\gamma(n)\right)+\rho R(w) \tag{7}$$

$$\text{s.t.} \quad \frac{\delta_{n,N}}{\Delta_N}z(n)=\frac{\delta_{n,N}}{\Delta_N}h_n^{\mathsf{T}}w,\quad n=1,2,\ldots,N \tag{8}$$

To see the equivalence, observe that if $\delta_{n,N}$ happens to be zero for index n, then it does not matter whether a constraint exists at that point in time or not because the corresponding loss term will disappear from the sum in (7). Next, we introduce the Lagrangian function [16]:

$$\mathcal{L}(w,z,\lambda)=\frac{1}{\Delta_N}\sum_{n=1}^{N}\Big(\delta_{n,N}Q\Big(z(n);\gamma(n)\Big)+\delta_{n,N}\lambda(n)z(n)\Big) + \rho\left(R(w)-\frac{1}{\rho\Delta_N}\sum_{n=1}^{N}\delta_{n,N}\lambda(n)h_n^{\mathsf{T}}w\right) \tag{9}$$

where $\{\lambda(n)\}$ are scalar Lagrange multipliers. Observe that we have as many Lagrange multipliers as the number of data samples and, therefore, their number increases continuously in the case of streaming data (which is the situation we are interested in). Next, we introduce the conjugate functions [16]:

$$Q^{\star}(x;\gamma) \;\triangleq\; \sup_z\left\{xz-Q(z;\gamma)\right\} \tag{10}$$

$$R^{\star}(x) \;\triangleq\; \sup_w\left\{x^{\mathsf{T}}w-R(w)\right\} \tag{11}$$

We can now express the dual function, denoted by $\mathcal{D}_N(\lambda)$, in terms of these conjugate function as follows:

$$\begin{aligned}\mathcal{D}_N(\lambda)&=\frac{1}{\Delta_N}\sum_{n=1}^{N}\delta_{n,N}\inf_{z(n)}\Big(Q\left(z(n);\gamma(n)\right)+\lambda(n)z(n)\Big)\\ &\quad+\rho\inf_w\left(R(w)-\frac{1}{\rho\Delta_N}\sum_{n=1}^{N}\delta_{n,N}\lambda(n)h_n^{\mathsf{T}}w\right)\\ &=-\frac{1}{\Delta_N}\sum_{n=1}^{N}\delta_{n,N}Q^{\star}\big(-\lambda(n);\gamma(n)\big)\\ &\quad-\rho R^{\star}\left(\frac{1}{\rho\Delta_N}\sum_{n=1}^{N}\delta_{n,N}\lambda(n)h_n\right)\end{aligned} \tag{12}$$

From the infimum operation on the regularization term, the primal variable w_N has to satisfy the following first-order optimality condition:

$$\frac{1}{\rho\Delta_N}\sum_{n=1}^{N}\delta_{n,N}\lambda(n)h_n\in\partial R(w_N) \tag{13}$$

where $\partial R(w)$ denotes the sub-differential of $R(\cdot)$ at point w. Now it is known that if a function $F(\cdot)$ is convex and closed, then it holds that [17]:

$$x\in\partial F(y)\Longleftrightarrow y\in\partial F^{\star}(x) \tag{14}$$

Applying this property to (13) and recalling that $R(w)$ is assumed to be strongly-convex, which implies its conjugate function, $R^{\star}(x)$, is continuously differentiable [17], we find that the primal variable w_N is given by the following expression in terms of the gradient vector of the conjugate regularization function:

$$w_N=\nabla_x R^{\star}\left(w_N'\right) \tag{15}$$

where we introduced the intermediate variable:

$$w_N' \;\triangleq\; \frac{1}{\rho\Delta_N}\sum_{n=1}^{N}\delta_{n,N}\lambda(n)h_n \tag{16}$$

TABLE I
TYPICAL CHOICES FOR THE REGULARIZATION TERM, ITS CONJUGATE FUNCTION AND GRADIENT VECTOR.

	$R(w)$	$R^\star(x)$	$\nabla_x R^\star(x)$
a)	$\frac{1}{2}\|w\|^2$	$\frac{1}{2}\|x\|^2$	x
b)	$\delta\|w\|_1+\frac{1}{2}\|w\|^2$	$\frac{1}{2}\|\mathcal{T}_\delta(x)\|^2$	$\mathcal{T}_\delta(x)$
c)	$\sum_{i=1}^M w(i)\log w(i)$	$\log\left(\sum_{i=1}^M e^{x(i)}\right)$	$\frac{e^x}{\sum_i e^{x(i)}}$

a) ℓ_2-regularization.
b) Elastic-net regularization. The entry-wise soft-threshold operator is defined as $[\mathcal{T}_\delta(w)]_i \triangleq (|w(i)|-\delta)_+\mathrm{sgn}(w(i))$ for the $i-$th entry.
c) Regularization based on KL divergence. Here, $w(i)$ represents the $i-$th entry of w and vector w belongs to the probability simplex.

In Table I we list some common choices for the regularization term, its conjugate function, and gradient vector. In this way, the dual function from (12) can be expressed as:

$$\mathcal{D}_N(\lambda) = -\frac{1}{\Delta_N}\sum_{n=1}^N \delta_{n,N} Q^\star\big(-\lambda(n);\gamma(n)\big) - \rho R^\star(w'_N) \quad (17)$$

B. Recursive constructions

We still need to determine the dual variables, $\{\lambda(n)\}$, which help identify the primal solution through (15) and (16). Before showing how to carry out this calculation, we observe first that expression (16) for the intermediate variable allows us to motivate recursive constructions for the primal variable. Revisiting the three scenarios we considered before:

(C1) (**Infinite-length window)**: In this case, we have

$$w'_N = \frac{N-1}{N}w'_{N-1} + \frac{1}{\rho N}\lambda(N)h_N \quad (18)$$

(C2) (**Exponential-weighting window**): In this case, we have

$$w'_N = \frac{\beta-\beta^N}{1-\beta^N}w'_{N-1} + \frac{1-\beta}{\rho(1-\beta^N)}\lambda(N)h_N \quad (19)$$

(C3) (**Finite-length sliding window**): We only consider the $N>L$ situation; the case $N\leq L$ can be handled similarly.

$$w'_N = w'_{N-1} - \frac{1}{\rho L}\lambda(N-L)h_{N-L} + \frac{1}{\rho L}\lambda(N)h_N \quad (20)$$

In all cases (18)–(20), we find that there is a mapping that transforms w'_{N-1} into w'_N. We denote this mapping generically by

$$w'_N = f_N(w'_{N-1}) + \alpha(N)\lambda(N)h_N \quad (21)$$

for some scalar $\alpha(N)$ and function $f_N(x)$ given by (the function $f_N(\cdot)$ is affine in these three examples):

$$f_N(x) = \begin{cases} \frac{N-1}{N}x, & \text{(infinite-length window)} \\ \frac{\beta-\beta^N}{1-\beta^N}x, & \text{(exponential window)} \\ x - \frac{1}{\rho L}\lambda(N-L)h_{N-L}, & \text{(sliding window)} \end{cases}$$

and

$$\alpha(N) = \begin{cases} \frac{1}{\rho N}, & \text{(infinite-length window)} \\ \frac{1-\beta}{\rho(1-\beta^N)}, & \text{(exponential window)} \\ \frac{1}{\rho L}, & \text{(sliding window)} \end{cases} \quad (22)$$

Observe that in the three cases considered above it holds that

$$\alpha(N) = \frac{1}{\rho\Delta_N} \quad (23)$$

C. Online algorithm

Observe from (17) that the dual function $\mathcal{D}_N(\lambda)$ depends on both N and λ, which creates a challenge for the development of an online algorithm. This is because the form of the dual function changes with N. Also, the number of dual variables increases with N. We therefore need an efficient method to seek the maximizer of the dual function. The main idea is as follows. When a new data point $\{\gamma(N), h_N\}$ streams in, we shall fix the previous Lagrange multipliers $\{\lambda(1),\lambda(2),\ldots,\lambda(N-1)\}$ at their existing values and then maximize $\mathcal{D}_N(\lambda)$ only with respect to $\lambda(N)$. It is important to emphasize that the motivation for this argument is somewhat different from traditional coordinate-ascent implementations. This is because the number of dual variables is now changing with time and, therefore, it is not possible to simply start from the solution of the last iteration. Instead, we extend the last solution into an enlarged vector that is one dimension higher and fix the leading entries of this longer vector to the dual variables from the last iteration. In this way, we can write the dual function (17) as

$$\begin{aligned}\mathcal{D}_N(\lambda) \overset{(21)}{=} &-\frac{1}{\Delta_N}Q^\star\Big(-\lambda(N);\gamma(N)\Big) + \text{const} \\ &- \rho R^\star\left(f_N(w'_{N-1}) + \alpha(N)\lambda(N)h_N\right)\end{aligned} \quad (24)$$

where the term "const" aggregates terms that are independent of $\lambda(N)$. By maximizing over $\lambda(N)$ we arrive at the proposed online dual coordinate-ascent (O-DCA) algorithm:

$$\begin{aligned}\lambda(N) &= \arg\min_\tau \frac{1}{\Delta_N}Q^\star\Big(-\tau;\gamma(N)\Big) \\ &\qquad + \rho R^\star\left(f_N(w'_{N-1}) + \tau\cdot\alpha(N)h_N\right) && \text{(25a)} \\ w'_N &= f_N(w'_{N-1}) + \alpha(N)\lambda(N)h_N && \text{(25b)} \\ w_N &= \nabla R^\star(w'_N) && \text{(25c)}\end{aligned}$$

Observe that the algorithm involves three steps at each iteration N, when a new data $\{\gamma(N), h_N\}$ streams in. First, the optimal $\lambda(N)$ is determined by solving (25a) Then, the intermediate estimate w'_N is determined, followed by the evaluation of w_N. In comparison with the stochastic DCA (S-DCA) implementation of [13], [14], three main differences stand out. First, at each iteration N, the proposed algorithm (25a)–(25c) is employing a time-varying normalization factor Δ_N, rather than a fixed N. This feature is critical for handling streaming data and to enable adaptation and tracking. Second, each data $\{\gamma(N), h_N\}$ is only used once, which is necessary for streaming data scenarios; multiple passes over the data are not practical in this case. And, third, more weighting is assigned to recent data than past data, which is important for scenarios with drifting minimizers.

In cases when the loss function $Q(\cdot)$ is non-differentiable, it is often helpful to smooth the output of O-DCA as follows:

$$\bar{w}_N \triangleq \frac{1}{S_N}\sum_{n=1}^N \kappa^{N-n}w_n, \quad \text{where } S_N \triangleq \sum_{n=1}^N \kappa^{N-n} \quad (26)$$

and the weight factor $\kappa \in [0,1]$. Computing $\bar{w}_N$ can be implemented efficiently, e.g., by using the same recursive method used before in (19) to find that.

D. Relation to Stochastic Primal Algrithms

The online DCA algorithm (25a)–(25c) that we just derived has strong connections with learning algorithms in the primal domain, especially when ℓ_2-regularization is employed, i.e., $R(w) = \frac{1}{2}\|w\|^2$. In this case, solving the argmin problem in (25a) requires that we determine a $\lambda(N)$ that satisfies the following first-order condition:

$$h_N^{\mathsf{T}} \nabla R_x^{\star}\Big(f_N(w'_{N-1}) + \alpha(N)\lambda(N)h_N\Big) \in \partial Q^{\star}\Big(-\lambda(N);\gamma(N)\Big) \tag{27}$$

Using property (14) and update step (25b), we conclude that:

$$\lambda(N) \in -\partial Q\Big(h_N^{\mathsf{T}} \nabla_x R^{\star}(w'_N);\gamma(N)\Big) \tag{28}$$

It follows that we can rewrite the update for the intermediate variable in the O-DCA algorithm (25a)–(25b) as follows:

$$w'_N = f_N(w'_{N-1}) - \alpha(N)\partial Q\left(h_N^{\mathsf{T}} \nabla R_x^{\star}(w'_N);\gamma(N)\right) h_N \tag{29}$$

For ℓ_2-regularization we have $\nabla_x R^{\star}(x) = x$ and therefore the above recursion, along with the last equality (25c), show that the proposed O-DCA algorithm reduces to the following insightful form (notice that w_N appears inside the sub-differential instead of w_{N-1}):

$$w_N^{\mathrm{ODCA}} = f_N(w_{N-1}^{\mathrm{ODCA}}) - \alpha(N)\partial Q\Big(h_N^{\mathsf{T}} w_N^{\mathrm{ODCA}};\gamma(N)\Big)h_N \tag{30}$$

This form reveals important connections with iterative algorithms in the primal domain. Indeed, note that if we were to use an online stochastic (sub-)gradient descent (SGD) method to solve (1) with ℓ_2- regularization, we would have obtained the following recursion:

$$w_N^{\mathrm{SGD}} = (1-\mu\rho)w_{N-1}^{\mathrm{SGD}} - \mu\partial Q\Big(h_N^{\mathsf{T}} w_{N-1}^{\mathrm{SGD}};\gamma(N)\Big)h_N \tag{31}$$

On the other hand, if we were to use an online stochastic proximal gradient (SPG) method to solve (1), again with ℓ_2- regularization, we would have obtained:

$$\begin{aligned} w_N^{\mathrm{SPG}} &= \mathrm{prox}_{\mu Q}\Big((1-\mu\rho)w_{N-1}^{\mathrm{SPG}}\Big) \\ &= (1-\mu\rho)w_{N-1}^{\mathrm{SPG}} - \mu\partial Q\Big(h_N^{\mathsf{T}} w_N^{\mathrm{SPG}};\gamma(N)\Big)h_N \end{aligned} \tag{32}$$

where in the last step we used the fact that $u = \mathrm{prox}_Q(x) \Longleftrightarrow x - u \in \partial Q(u)$. Comparing (30), (31), and (32) under ℓ_2-regularization, we observe that although O-DCA was formulated in the dual domain, it can still be viewed as one form of a proximal implementation with the variable w_N^{OCDA} appearing on the right-hand side of (30) inside the sub-differential term, as happens with w_N^{SPG} in (32).

III. Specific Loss functions

We illustrate the above connections more explicitly by considering two important cases: least-mean-squares error designs and support vector machines. In both cases, for simplicity, we continue to employ ℓ_2-regularization, $R(w) = \frac{1}{2}\|w\|^2$.

A. Least-Mean-Squares Learning

In this case we have

$$Q\Big(z(N);\gamma(N)\Big) = \frac{1}{2}\big(\gamma(N) - z(N)\big)^2 \tag{33}$$

$$Q^{\star}\Big(\lambda(N);\gamma(N)\Big) = \frac{1}{2}\lambda^2(N) + \gamma(N)\lambda(N) \tag{34}$$

Therefore, assuming an infinite-length window, we need to solve the following optimization problem to find $\lambda(N)$:

$$\min_{\tau} \frac{1}{N}\left(\frac{1}{2}\tau^2 - \gamma(N)\tau\right) + \frac{\rho}{2}\left\|\frac{N-1}{N}w_{N-1} + \frac{\tau h_N}{\rho N}\right\|^2 \tag{35}$$

Setting the derivative relative to τ equal to zero at $\tau = \lambda(N)$ we find that the O-DCA algorithm (25a)–(25c) reduces to:

$$\lambda(N) = \left(1 + \frac{\|h_N\|^2}{\rho N}\right)^{-1}\left(\gamma(N) - \frac{N-1}{N}h_N^{\mathsf{T}} w_{N-1}\right) \tag{36a}$$

$$w_N = \frac{N-1}{N}w_{N-1} + \frac{\lambda(N)h_N}{\rho N} \tag{36b}$$

Assuming N is large enough, we can merge the two equations into:

$$w_N = w_{N-1} + \frac{1}{\rho N}h_N\Big(\gamma(N) - h_N^{\mathsf{T}} w_{N-1}\Big) \tag{37}$$

The O-DCA implementation (37) approaches a leaky-LMS implementation with a decaying step-size of the form $\mu_N = 1/\rho N$, namely,

$$w_N = (1-\rho\mu_N)w_{N-1} + \mu_N h_N(\gamma_N - h_N^{\mathsf{T}} w_{N-1}) \tag{38}$$

B. Support-Vector Machines

In this case, we have (where $\gamma(N) = \pm 1$):

$$Q\Big(z(N);\gamma(N)\Big) = \max\{0, 1-\gamma(N)z(N)\} \tag{39}$$

with conjugate function (in compact form):

$$Q^{\star}\big(\lambda(N);\gamma(N)\big) = \begin{cases} \gamma(N)\lambda(N), & \text{if } \gamma(N)\lambda(N) \in [-1,0] \\ +\infty, & \text{otherwise} \end{cases} \tag{40}$$

In this example, we consider the exponential weighting window so that the dual variable is found by solving:

$$\begin{aligned} \lambda(N) = \arg\min_{\tau} & \left\{-\frac{\gamma(N)\tau}{\Delta_N} + \frac{\rho}{2}\left\|\frac{\beta-\beta^N}{1-\beta^N}w_{N-1} + \frac{\tau h_N}{\rho\Delta_N}\right\|^2\right\} \\ & \text{subject to} \quad \gamma(N)\tau \in [0,1] \end{aligned} \tag{41}$$

This optimization problem involves a truncated parabola function as a cost objective. The minimizer of the quadratic cost occurs at

$$\lambda(N) = \frac{\rho\Delta_N}{\|h_N\|^2}\left(\gamma(N) - \frac{\beta-\beta^N}{1-\beta^N}h_N^{\mathsf{T}} w_{N-1}\right) \tag{42}$$

We still need to adjust this value, by means of a projection operation, in order to meet the constraint. To simplify the projection, we multiply both sides of the above relation by $\gamma(N)$ and use $\gamma^2(N) = 1$ to find that the O-DCA algorithm (25a)–(25c) reduces to:

$$\gamma(N)\lambda(N) = \frac{\rho\Delta_N}{\|h_N\|^2}\left(1 - \frac{\beta-\beta^N}{1-\beta^N}\gamma(N)h_N^{\mathsf{T}}w_{N-1}\right) \tag{43}$$

$$w_N = \frac{\beta-\beta^N}{1-\beta^N}w_{N-1} + \frac{1-\beta}{\rho(1-\beta^N)}\gamma(N)h_N\prod_{[0,1]}\left[\gamma(N)\lambda(N)\right]$$

where $\Pi_{[0,1]}(a)$ projects the real number a into the interval $[0,1]$. If we let $\beta = 1-\mu\rho$ and assume a small enough μ so that the value of β is close to one, the above two equations can be merged into the following format when N is large enough:

$$w_N^{\mathrm{ODCA}} = (1-\mu\rho)w_{N-1}^{\mathrm{ODCA}} + \mu\gamma(N)h_N \cdot \prod_{[0,1]}\left[\frac{1}{\mu\|h_N\|^2}\left(1-\gamma(N)h_N^{\mathsf{T}}w_{N-1}^{\mathrm{ODCA}}\right)\right] \tag{44}$$

For comparison purposes, we list the stochastic subgradient solution for SVM here [8], [18]:

$$w_N^{\mathrm{SGD}} = (1-\mu\rho)w_{N-1}^{\mathrm{SGD}} + \mu\gamma(N)h_N \cdot \mathbb{I}[1-\gamma(N)h_N^{\mathsf{T}}w_{N-1}^{\mathrm{SGD}} \geq 0] \tag{45}$$

where $\mathbb{I}[\cdot]$ is the indicator function, which is equal to one when the argument is true and zero otherwise. Comparing (44) with (45), it is clear that O-DCA replaces the indicator function (which involves a sudden jump from 0 to 1) by a smoothed linear transition from 0 to 1 with slope proportional to $1/\mu$.

IV. Simulations

We illustrate the performance of O-DCA using the hinge loss (39) applied to two datasets. The test data is obtained from the LIBSVM website[1]. We first use the Adult dataset with 11,220 training data and 21,341 testing data in 123 feature dimensions. In the figure, we compare our algorithm with the S-DCA method from [13], the stochastic sub-gradient method from [18], and with LIBSVM [19]. To perform a fair comparison, we compare the performance based upon the number of iterations (except for LIBSVM). The parameter setting is as follows. All algorithms use $\rho = 0.001$. For O-DCA, we choose $\beta = 0.99995$ and the step-size for the sub-gradient implementation is 0.05, so that it satisfies $\beta = 1-\mu\rho$. Since S-DCA is not designed for online learning, we feed S-DCA with one-fifth of the training data, and run over 5 epochs so that the total number of iterations will match.

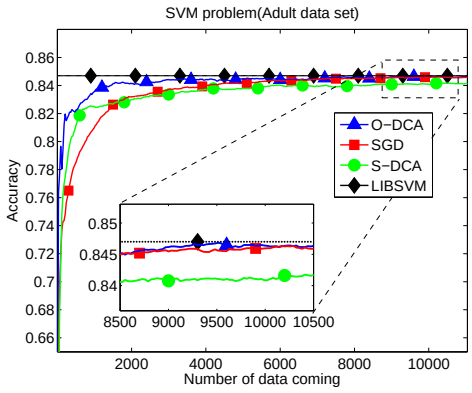

Fig. 1. Comparison of the performance of O-DCA, S-DCA, SGD, and LIBSVM on the Adult dataset in terms of iterations.

[1]http://www.csie.ntu.edu.tw/∼cjlin/libsvmtools/datasets

The second dataset is the Reuters Corpus Volume I (RCV1) data with 20242 training data and 253843 testing data consisting of 47236 feature dimensions. Similarly, we set $\rho = 10^{-4}$, $\beta = 0.99998$ for O-DCA, $\mu = 0.2$ for sub-gradient so that $\beta = 1-\mu\rho$, and epochs $= 5$ for S-DCA.

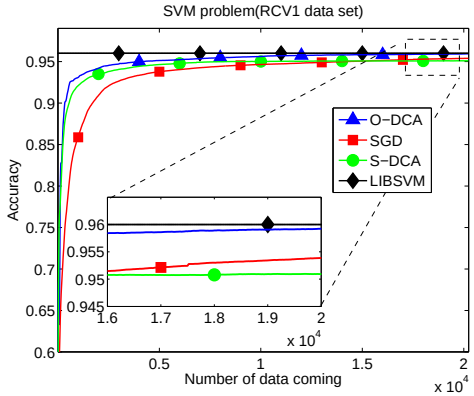

Fig. 2. Comparison of the performance of O-DCA, S-DCA, SGD, and LIBSVM on the RCV dataset in terms of iterations.

References

[1] A. H. Sayed, *Adaptive Filters*, John Wiley & Sons, 2008.

[2] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.

[3] C.-J. Hsieh, K.-W. Chang, C.-J. Lin, S. S. Keerthi, and S. Sundararajan, "A dual coordinate descent method for large-scale linear SVM," in *Proceedings of the International Conference on Machine Learning*, Helsinki, Finland, 2008, ACM, pp. 408–415.

[4] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning*, Springer, 2009.

[5] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, 2014.

[6] H. Robbins and S. Monro, "A stochastic approximation method," *The Annals of Mathematical Statistics*, pp. 400–407, 1951.

[7] T. Zhang, "Solving large-scale linear prediction problems using stochastic gradient descent algorithms," in *Proceedings of the International Conference on Machine Learning*, Banff, Canada, 2004, pp. 116–123.

[8] S. Shalev-Shwartz, Y. Singer, N. Srebro, and A. Cotter, "Pegasos: Primal estimated sub-gradient solver for SVM," *Mathematical Programming*, vol. 127, no. 1, pp. 3–30, 2011.

[9] B. T. Polyak, *Introduction to Optimization*, Optimization Software New York, 1987.

[10] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*, Athena Scientific, Singapore, 1997.

[11] Z.-Q. Luo and P. Tseng, "On the convergence of the coordinate descent method for convex differentiable minimization," *Journal of Optimization Theory and Applications*, vol. 72, no. 1, pp. 7–35, 1992.

[12] H.-F. Yu, F.-L. Huang, and C.-J. Lin, "Dual coordinate descent methods for logistic regression and maximum entropy models," *Machine Learning*, vol. 85, no. 1-2, pp. 41–75, 2011.

[13] S. S.-Shwartz and T. Zhang, "Stochastic dual coordinate ascent methods for regularized loss," *Journal of Machine Learning Research*, vol. 14, no. 1, pp. 567–599, 2013.

[14] S. S.-Shwartz and T. Zhang, "Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization," *Mathematical Programming*, pp. 1–41, 2014.

[15] D. Angelosante, J.-A. Bazerque, and G. B. Giannakis, "Online adaptive estimation of sparse signals: Where RLS meets the ℓ_1-norm," *IEEE Transactions on Signal Processing*, vol. 58, no. 7, pp. 3436–3447, 2010.

[16] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.

[17] D. Bertsekas and A. Nedic, *Convex Analysis and Optimization*, Athena Scientific, 2003.

[18] B. Ying and A. H. Sayed, "Performance limits of online stochastic sub-gradient learning," *available as arXiv:1511.07902*, Nov. 2015.

[19] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011.